CAS Severn
Intelligent Answers.

Expert Video Series: Illustrated Transcript

Creating a Spectrum Protect Cloud Container Pool on Amazon S3

YouTube

Click here to see the video: http://bit.ly/2BXCJWB

1978 40 2018
CAS Severn

Introduction: Hello, my name is Joseph King I'm the Chief Technology Officer at CAS Severn. Today's topic is how to create spectrum protect container pulls on Amazon S3 storage using local cache.

00:10 Spectrum Protect cloud container pools have a long history. It started back in third quarter of 2015 where they were introduced with Swift object storage, and there's been improvements every version, which is every quarter since then, to where we are today in version 8.1.3. with tearing to the cloud.

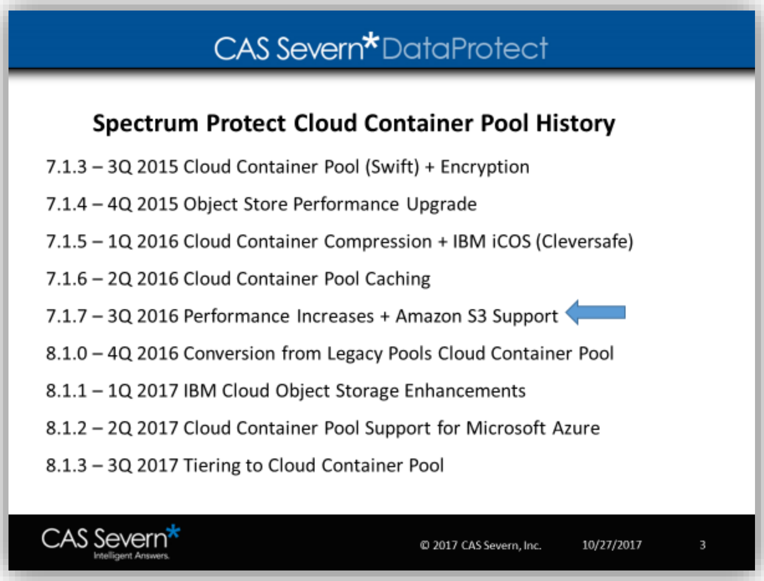

CAS Severn*DataProtect

Spectrum Protect Cloud Container Pool History

7.1.3 – 3Q 2015 Cloud Container Pool (Swift) + Encryption

7.1.4 – 4Q 2015 Object Store Performance Upgrade

7.1.5 – 1Q 2016 Cloud Container Compression + IBM iCOS (Cleversafe)

7.1.6 – 2Q 2016 Cloud Container Pool Caching

7.1.7 – 3Q 2016 Performance Increases + Amazon S3 Support

8.1.0 – 4Q 2016 Conversion from Legacy Pools Cloud Container Pool

8.1.1 – 1Q 2017 IBM Cloud Object Storage Enhancements

8.1.2 – 2Q 2017 Cloud Container Pool Support for Microsoft Azure

8.1.3 – 3Q 2017 Tiering to Cloud Container Pool

CAS Severn* Intelligent Answers. © 2017 CAS Severn, Inc. 10/27/2017 3

00:33 Today what we're going to talk about on the integration with Amazon S3 started in version 7.1.7 and third quarter of 2016. So if you're at 7.1.7 or beyond the topic that we're going to talk about you're going to be able to configure on your Spectrum Protect system.

00:53 There's been upgrades to the operation center in order to make this much easier to deploy, and you're going to see that in just a few steps we could be sending data to Amazons S3.

01:08 There are three steps in order to create an Amazon S3 cloud container pool. First, we must do some work on Amazon. Specifically, we have to create a group in a user that is going to hold the S3 storage that we're going to connect to spectrum protect.

01:21 Then we need to create a cache directory on the Spectrum Protect server. Lastly, we need to create the cloud container pool itself. Creating an Amazon user or group is very straightforward.

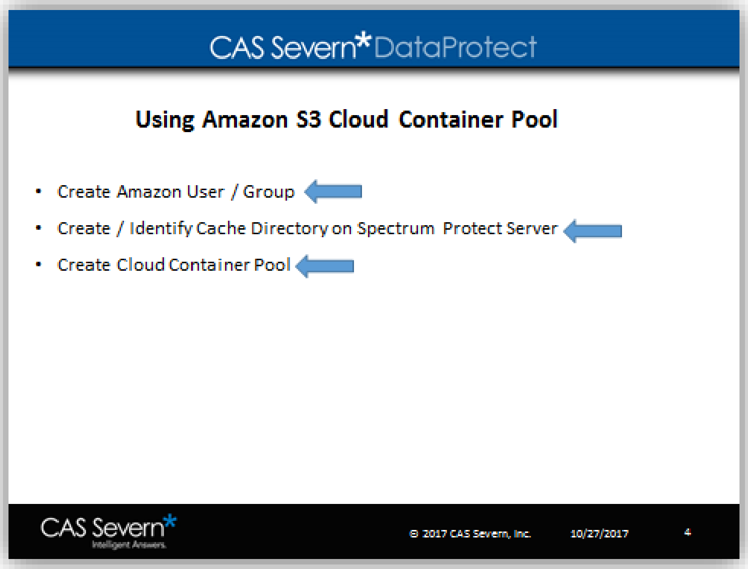

CAS Severn*DataProtect

Using Amazon S3 Cloud Container Pool

- Create Amazon User / Group
- Create / Identify Cache Directory on Spectrum Protect Server
- Create Cloud Container Pool

© 2017 CAS Severn, Inc. 10/27/2017 4

01:37 So first, we're going to create an Amazon group and we're going to give it full permissions to S3. Then we're going to create an Amazon user with programmatic access. We're then going to assign the user to a group, and we're going to download the access key of the secret key.

01:48 You might have individual security requirements for your company that you should adhere to creating your S3 buckets, and your users and groups. So, we're going to go through some defaults here that'll work for most cases. But again, follow your organization's security practice.

02:08 Starting in Amazon open your IAM tool and go to groups. You're going to create a new group, and for our demo we're going to use SP test for the group name. Next, pick a policy. There are some default policies that are integrated in to Amazon AWS. The one that we're going to use today is Amazon S3 full access.

02:45 Lastly, we are going to review before creating the group review your settings. Make sure the group name set and the policy set to what we like. Once you are happy with the settings create the group.

02:58 Next, we're going to create a user. So in IAM go to the users area, add user, and we're going to use Spectrum Protect test user or SP test user here and programmatic access.

03:17 Next, we're going to work with permissions. Here, because we've already created the group we can attach that group we just created right to that user.

03:27 Review that we chose the right group and create user. Lastly, because we chose access we have to download the access key and the secret access key. The easiest way to do that is download the CSV file.

03:52 I downloaded that to my desktop and I'm going to open it up here in notepad and grab the secret key ID and the secret key itself.

04:09 Now we're going to switch back and talk about what we're going to do within Spectrum Pprotect.

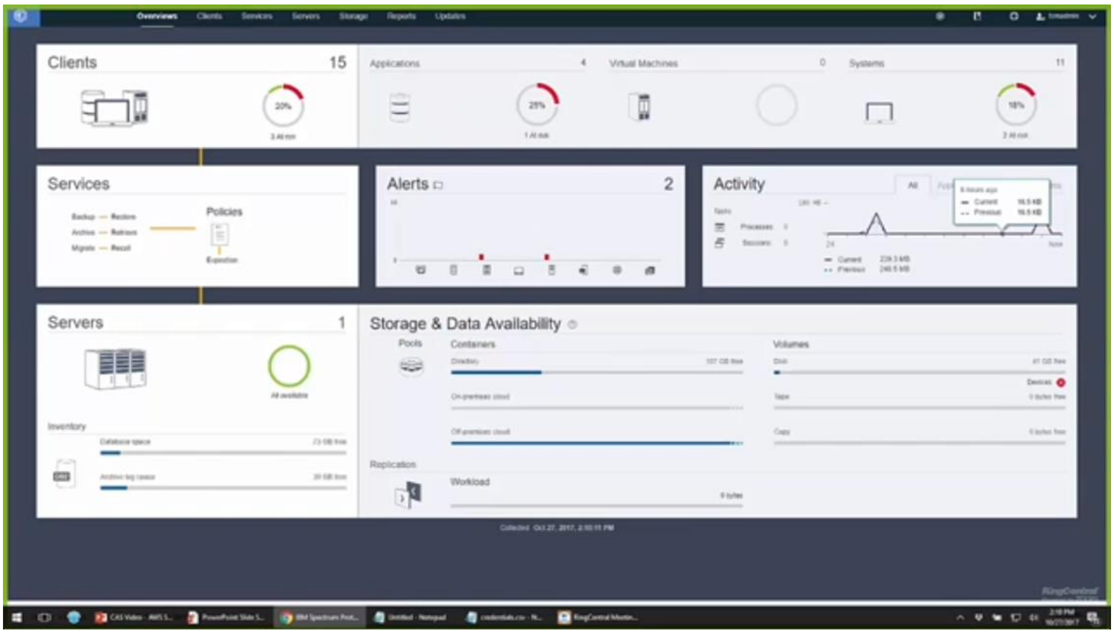

So, we logged in an operation center as an administrative user and go to storage and storage pools. You can see we already have a cloud container pool created here called Amazon S3 pool. We can see that we have some savings there due to deduplication and compression.

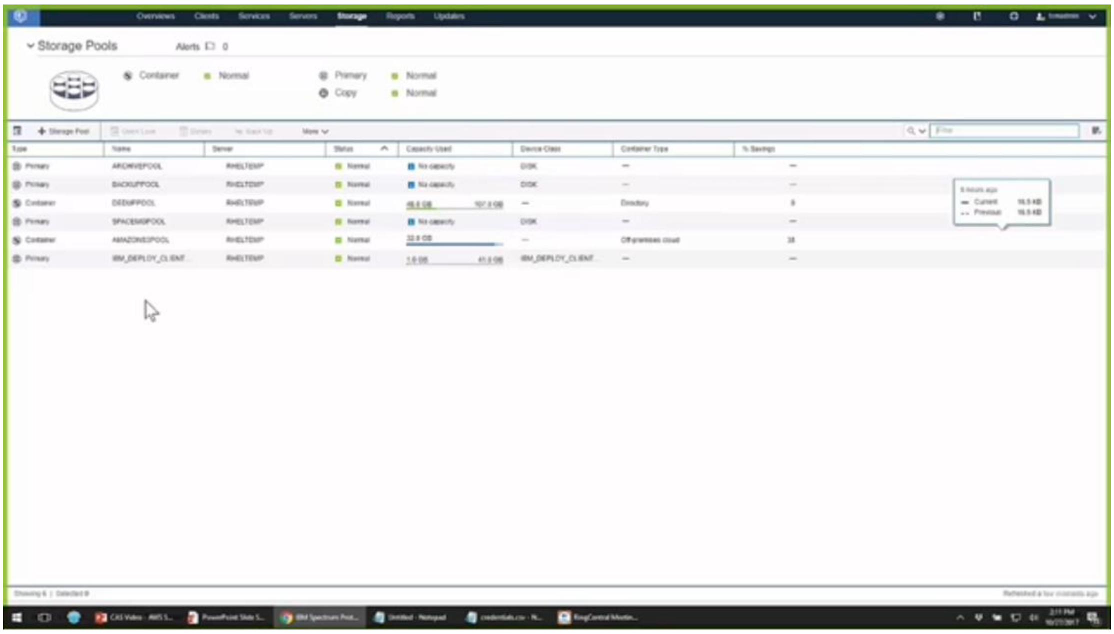
Storage Pools Alerts 0

Container ■ Normal
Primary ■ Normal
Copy ■ Normal

Type	Name	Server	Status	Capacity Used	Device Class	Container Type	% Savings
Primary	ARCHIVEPOOL	RHELTEMP	Normal	No capacity	DISK	—	—
Primary	BACKUPPOOL	RHELTEMP	Normal	No capacity	DISK	—	—
Container	DEDUPPOOL	RHELTEMP	Normal	48.8 GB 107.8 GB	—	Directory	8
Primary	SPACEMGPOOL	RHELTEMP	Normal	No capacity	DISK	—	—
Container	AMAZONS3POOL	RHELTEMP	Normal	32.8 GB	—	Off-premises cloud	38
Primary	IBM_DEPLOY_CLIENT	RHELTEMP	Normal	1.6 GB 41.8 GB	IBM_DEPLOY_CLIENT	—	—

04:45 Cloud container pools can be deduplicated, compressed and encrypted just like directory container pools can. So let's create a new one.

04:58 Choose the plus button there,   give it a name make sure that we're connecting to the correct Spectrum Protect server. Choose off-premise cloud from your options cloud type Amazon S3 API. And here's where you're going to enter the access key ID and the secret access key that you downloaded in the credential CSV file in the previous step.

05:42 This is where to choose the cache directory. The cache directory is where the data is going to be initially transferred to when it's ingested off the clients, and in that cache directory the data will be deduplicated, compressed or encrypted, or all three depending on your settings.

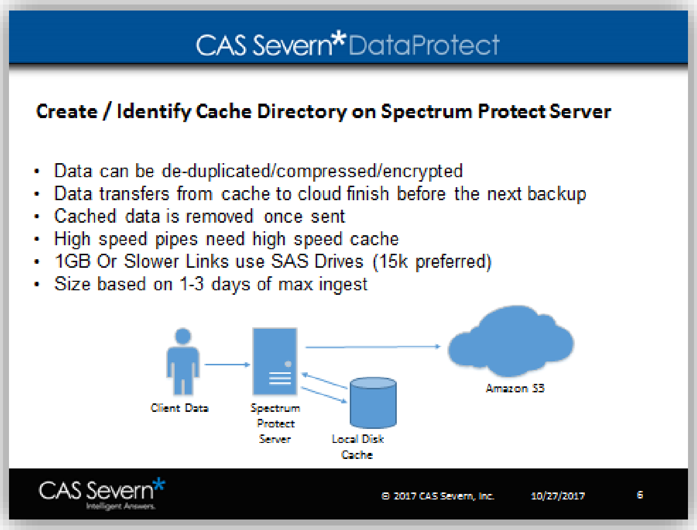

CAS Severn*DataProtect

Create / Identify Cache Directory on Spectrum Protect Server

- Data can be de-duplicated/compressed/encrypted
- Data transfers from cache to cloud finish before the next backup
- Cached data is removed once sent
- High speed pipes need high speed cache
- 1GB Or Slower Links use SAS Drives (15k preferred)
- Size based on 1-3 days of max ingest

CAS Severn* Intelligent Answers.   © 2017 CAS Severn, Inc.   10/27/2017   6

06:09 Then this data transfer to cache is going to help speed the data ingest to the Spectrum Protect server. So first the data is going to come from the client and be stored on local cache. Then at the end of the backup cycle, that data then is going to be sent up to the cloud container pool on Amazon S3.

06:32 How you build a local cache is highly dependent on the speed of your network to Amazon. If you have a dedicated network to Amazon using the direct connect, then you're going to need a high-speed pipe.

06:47 At that point, your cache also needs to be high-speed. So if you have a ten-gig connection to Amazon S3, solid-state drives for the cache make a lot of sense. If you have a one gigabyte or gigabit or slower link, then the use of SAS drives is fine. While 15,000 rpm drives are preferred, 10,000 worked just fine. When it comes to sizing, size based off at least your max daily ingest. But an easy way of dealing with this is just take one to three days of your ingest, post deduplication and compression, and use that to size your cache.

07:39 I've already created a temporary directory to use for this demonstration, and I just put it in temp for this. You're going to have to put it in a file system that makes sense for your workload. So, in temp storage pool, add a storage pool, so this is going to create a bucket on Amazon S3, and then it's going to create a cloud container pool.

08:09 This takes a minute or two, and once it creates the bucket and it creates the pool, you should get green checkboxes all around a big one at the top says succeeded, and at that point the cloud container pool is created.

08:37 To use the cloud container pool, you're going to have to modify one of your management policies in order to use it. In this case, we already have a policy setup and standard that's going to that pool that we created earlier.

08:55 Some details in creating a cloud container pool. It can be done in the command line. Though, as you've seen, doing it in operation center is very easy and the bucket names must be unique. We recommend going with the default bucket names that are created and operation center and those buckets will be maintained then by Spectrum Protect.

09:15 Let's take a look at what those buckets look like on Amazon. Now we're going to log in to Amazon and go into the S3 area. As you can see, the bucket that was created by Spectrum Protect is there.

09:29 You should see what region it was created in which is the default region that I'm in and the data it was created. Spectrum protect will take care of all the bucket management as you can see. Creating an S3 cloud container pool is easy, just create a group and user, create a cache directory and then we created a cloud container pool.

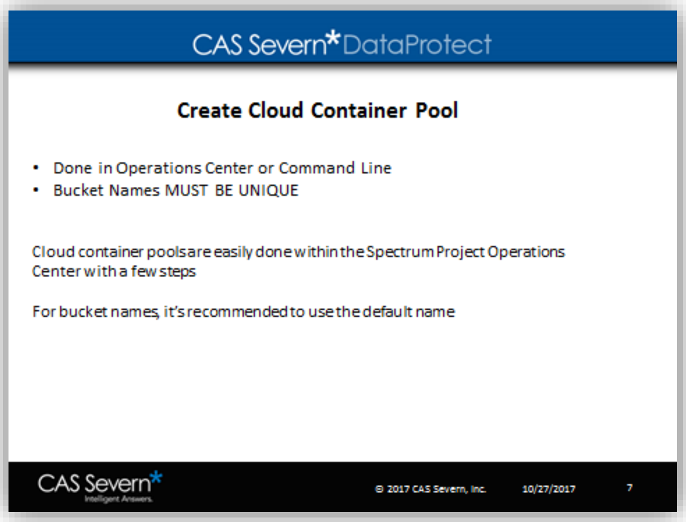

CAS Severn* DataProtect

Create Cloud Container Pool

- Done in Operations Center or Command Line
- Bucket Names MUST BE UNIQUE

Cloud container poolsare easily done within the Spectrum Project Operations Center with a few steps

For bucket names, it's recommended to use the default name

CAS Severn* Intelligent Answers. © 2017 CAS Severn, Inc. 10/27/2017 7

09:49 With that you can send your data to the cloud. Creating cloud container pulls on Amazon S3 is easy. Again my name is Joseph King chief technology officer of CAS Severn feel free to contact us if you have any questions.

To schedule a 1:1 with Joe and his technical team, contact us at:

sales@cassevern.com

800-252-4715